Journal of Scientific Computing manuscript No.
(will be inserted by the editor)

Nonuniformly weighted Schwarz smoothers for spectral element multigrid

Jörg Stiller

Received: date / Accepted: date

arXiv:1512.02390v1 [cs.NA] 8 Dec 2015

Abstract A hybrid Schwarz/multigrid method for spectral element solvers to the Poisson equation in $\mathbb{R}^2$ is presented. It is based on an additive Schwarz method studied by J. Lottes and P. Fischer (J. Sci. Comput. 24:45–78, 2005), which makes use of the inverse counting matrix for weighting overlapping Schwarz updates. It is shown that by introducing a nonuniform weighting function resembling a smoothed jump, the logarithmic multigrid convergence rate is improved by a factor of 1.5 to 3, leading to a corresponding reduction of the iteration count in comparison to the original method. Further, the influence of the overlap width, smoothing strategies, additive versus multiplicative Schwarz methods, and Krylov acceleration on robustness and efficiency is investigated.

Keywords Multigrid method · Schwarz methods · spectral element method · p-version finite element method

1 Introduction

High-order finite element methods (FEM) enjoy an increasing interest in computational science and engineering. They include hp-FEM, spectral element methods (SEM) as well as discontinuous Galerkin methods [5, 16]. The motive that drives the development of high-order methods lies in their potential to deliver accuracy with lower cost in comparison to the first and second order methods used in common simulation tools [29]. However, realizing this advantage in practice is a formidable task. Along with curvilinear mesh generation,

J. Stiller
Technische Universität Dresden, Institute of Fluid Mechanics (ISM) and
Center of Advancing Electronics Dresden (cfAED)
01062 Dresden, Germany
Tel.: +49-351-46338328
Fax: +49-351-46335246
E-mail: joerg.stiller@tu-dresden.de

the provision of efficient solvers for the resulting algebraic equation systems remains the main challenge.

Projection methods for incompressible flow, or implicit discretization of diffusion terms lead to a sequence of linear elliptical problems which are related or equivalent to the Poisson equation or, more generally, the Helmholtz equation [9]. Fast solvers for such equations are therefore a crucial ingredient of competitive high-order methods and, hence, have been in focus of research for almost 30 years [1, 3, 7, 10, 11, 13, 15, 17–19, 21, 22, 24–26]. For Helmholtz or Poisson problems discretized on regular meshes, efficient multigrid (MG) techniques have been developed recently. Lottes and Fischer [19] proposed additive Schwarz smoothers based on extended element domains, which attain residual reductions of approximately 0.2 within one sweep. They found that weighting the overlapping Schwarz updates by the inverse of the counting matrix, which corresponds to taking the arithmetic mean, plays a crucial role in obtaining multigrid-like iteration counts. A detailed analysis of the method was given in [18]. Janssen and Kanschat [13] presented a similar multigrid approach for the p-finite element method on locally refined Cartesian meshes. They used a multiplicative Schwarz smoother on element domains which possess only a minimal overlap confined to the element boundaries or, respectively, the corresponding DOF. Haupt et al. [10] developed a p-multigrid method based on static condensation which, apart from pre- and post-processing, reaches linear complexity. The proposed block smoother can be classified as an additive Schwarz method using a monotonic increasing shape function for weighting the overlapping updates. Using this smoother the multigrid method attained convergence rates of about 0.02 combined with a run-time efficiency that comes close to fast direct finite difference solvers. Unfortunately, despite some progress in factorizing the condensed operators [12], the approach is hard to generalize to three dimensions. Yet it inspired us, to extend the idea of nonuniform weighting to the full, "uncondensed" problem and thus led to the present work. Our primary goal is to show how nonuniform weighting can be used to boost the hybrid Schwarz-multigrid technique developed in [7] and [19]. Further we investigate the influence of the overlap width, smoothing strategies, additive versus multiplicative Schwarz methods, and Krylov acceleration on robustness and efficiency.

The remainder of the paper is organized as follows: Section 2 provides a brief description of the spectral element discretization. Section 3 presents the solution techniques, namely the weighted additive and multiplicative Schwarz methods, the p-multigrid method and the inexact multigrid-preconditioned conjugate gradient method. Section 4 proceeds with the discussion of numerical experiments for assessing the solution methods. Finally, Section 5 concludes the paper.

2 Discretization

As the model problem we consider the Poisson equation

$$-\nabla^2 u = f \tag{1}$$

in the rectangular domain $\Omega = [0, l_x] \times [0, l_y]$ with periodic boundaries. For discretization Ω is decomposed into in $n_\mathrm{E} = n_x \times n_y$ rectangular elements Ω^{mn} with dimensions $\Delta x = l_x/n_x$ and $\Delta y = l_y/n_y$. In each element the solution is approximated as

$$u(x,y)|_{\Omega^{mn}} \simeq \sum_{i,j=0}^{p} u_{ij}^{mn}\, \ell_i\big(\xi^m(x)\big)\, \ell_j\big(\eta^n(x)\big) \tag{2}$$

where ℓ_i are the Lagrange polynomials to the Gauss-Lobatto-Lengendre (GLL) points $\{\xi_i\}_{i=0}^p$ in the one-dimensional standard region $\hat{\Omega} = [-1, 1]$, $\xi^m(x)$ and $\eta^n(y)$ the mapping of coordinates from Ω^{mn} to $\hat{\Omega}$ [5, 16]. Concatenation of the element coefficients $\boldsymbol{u}^{mn} = [u_{ij}^{mn}]$ and enforcing continuity for shared vertices and edges yields the unique global coefficients $\boldsymbol{u}$ [see, e.g. 5, pp. 191–194]. Application of the Galerkin spectral element method leads to the discrete equations

$$\boldsymbol{A}\boldsymbol{u} = \boldsymbol{f}\,. \tag{3}$$

As a consequence of the tensor product ansatz (2) and the Cartesian mesh, the global system matrix in Eq. (3) assumes the tensor product form

$$\boldsymbol{A} = \boldsymbol{M}_y \otimes \boldsymbol{L}_x + \boldsymbol{L}_y \otimes \boldsymbol{M}_x\,, \tag{4}$$

where $\boldsymbol{M}_*$ and $\boldsymbol{L}_*$ represent the one-dimensional mass and stiffness matrices for directions $* = x, y$, respectively. The detailed structure of these operators and underlying spectral element techniques are well described in literature [5, 16, 19] and therefore deliberately skipped here.

3 Solution methods

For solving Eq. (3) we consider polynomial multigrid (MG) and multigrid-preconditioned conjugate gradients (MGCG). Both approaches rely on Schwarz methods for smoothing. We first present the Schwarz methods and then sketch MG and MGCG.

3.1 Schwarz methods

Schwarz methods are iterative domain decomposition techniques which improve the approximate solution by parallel or sequential subdomain solves, leading to additive or multiplicative methods, respectively. Here we follow Lottes and Fischer [19] who used extended element regions as the subdomains.

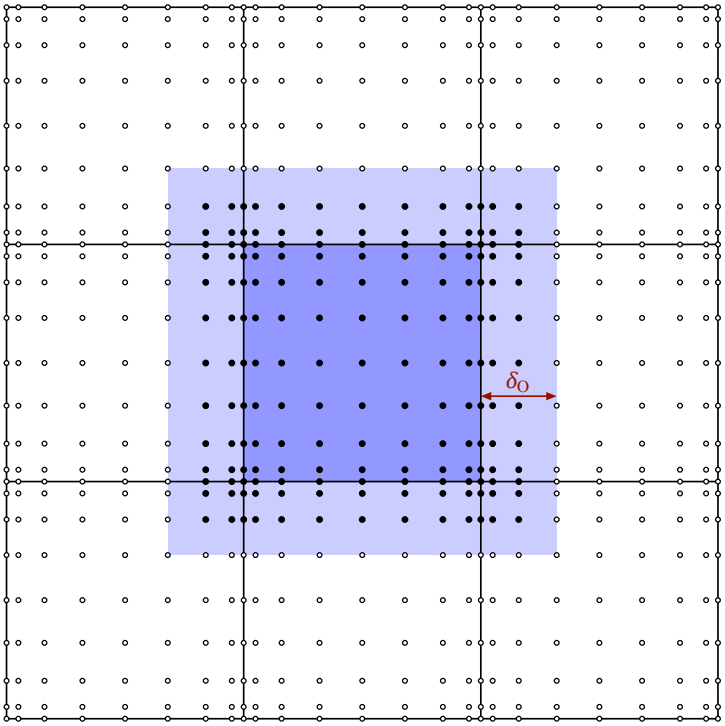

Fig. 1: Example of a subdomain used with the Schwarz method. The shaded area represents Ω_s and the dark region in its center the corresponding element. The circles are the GLL nodes for polynomial order $p = 8$. Filled circles indicate the nodes that are solved for and updated; δ_O is the overlap width.

Figure 1 illustrates how the subdomain $\Omega_{s(mn)}$ results from the corresponding element domain Ω^{mn} by attaching a rectangular strip matching the overlap width δ_O. As consequence, Ω_s adopts n_O layers of additional nodes from the neighbor elements. Note, however, that we exclude the outer layer of nodes located on $\partial\Omega_s$. For definiteness we define the overlap width in terms of n_O and the GLL points,

$$\delta_\mathrm{O} = \xi_{n_\mathrm{o}+1} + 1\,. \tag{5}$$

To derive a local correction to some approximate solution $\tilde{\boldsymbol{u}}$ we first convert Eq. (3) into the equivalent residual form

$$\boldsymbol{A}\Delta\boldsymbol{u} = \boldsymbol{f} - \boldsymbol{A}\tilde{\boldsymbol{u}} = \tilde{\boldsymbol{r}}\,, \tag{6}$$

where $\Delta\boldsymbol{u} = \tilde{\boldsymbol{u}} - \boldsymbol{u}$. Further we introduce the restriction operator $\boldsymbol{R}_s$ such that $\boldsymbol{u}_s = \boldsymbol{R}_s\boldsymbol{u}$ gives the coefficients associated with Ω_s. Conversely, the transposed restriction operator globalizes any local coefficients by adding zeros for exterior nodes. With these prerequisites the correction contributed by Ω_s is defined as

Algorithm 1 Multiplicative Schwarz method

function MSCHWARZ($\boldsymbol{u}, \boldsymbol{f}, n_\mathrm{I}$)
  for $i = 1, n_\mathrm{I}$ **do**
    for $e = 1, n_\mathrm{E}$ **do**
      $s \leftarrow \begin{cases} e & i \text{ odd} \\ n_\mathrm{E} + 1 - e & i \text{ even} \end{cases}$
      $\boldsymbol{r} \leftarrow \boldsymbol{f} - \boldsymbol{A}\boldsymbol{u}$
      $\boldsymbol{u} \leftarrow \boldsymbol{u} + \boldsymbol{R}_s^t \boldsymbol{A}_{ss}^{-1} \boldsymbol{R}_s \boldsymbol{r}$
    end for
  end for
  return $\boldsymbol{u}$
end function

Algorithm 2 Weighted additive Schwarz method

function WSCHWARZ($\boldsymbol{u}, \boldsymbol{f}, n_\mathrm{I}$)
  for $i = 1, n_\mathrm{I}$ **do**
    $\boldsymbol{r} \leftarrow \boldsymbol{f} - \boldsymbol{A}\boldsymbol{u}$
    $\boldsymbol{u} \leftarrow \boldsymbol{u} + \sum_{s=1}^{n_\mathrm{E}} \boldsymbol{R}_s^t \boldsymbol{w} \boldsymbol{A}_{ss}^{-1} \boldsymbol{R}_s \boldsymbol{r}$
  end for
  return $\boldsymbol{u}$
end function

the solution of the subproblem

$$\boldsymbol{A}_{ss} \Delta \boldsymbol{u}_s = \boldsymbol{r}_s \,, \tag{7}$$

where $\boldsymbol{A}_{ss} = \boldsymbol{R}_s \boldsymbol{A} \boldsymbol{R}_s^t$ represents the restricted system matrix and $\boldsymbol{r}_s = \boldsymbol{R}_s \tilde{\boldsymbol{r}}$ the restricted residual. Due to the rectangular shape of the subdomain, $\boldsymbol{A}_{ss}$ inherits the tensor product structure of $\boldsymbol{A}$ and can thus be inverted using the fast diagonalization technique developed by Lynch et al. [20] and adopted for SEM e.g. in [4, 19]. The inverse is the product of tensor product matrix to a diagonal matrix and another tensor product matrix such that the solution, $\Delta \boldsymbol{u}_s = \boldsymbol{A}_{ss}^{-1} \boldsymbol{r}_s$, can be evaluated with just $\Theta\big(2(p+1+2n_\mathrm{O})^3\big)$ operations.

There exist several options for combining the local solutions. We consider a weighted version of the additive Schwarz method and the multiplicative Schwarz method. The multiplicative Schwarz method solves the subproblems (7) consecutively while continually updating the residual. One multiplicative Schwarz iteration, in general, corresponds to the application of a nonsymmetric linear operator, albeit $\boldsymbol{A}$ is symmetric. However, for an even number of steps, the method is symmetrized by reversing the order of subdomains in each step, which finally leads to Algorithm 1.

The weighted additive Schwarz method determines all local corrections independently and computes the global correction as a linear combination of these results, i.e.

$$\Delta \boldsymbol{u} \simeq \sum_s \boldsymbol{R}_s^t (\boldsymbol{w} \Delta \boldsymbol{u}_s) \,, \tag{8}$$

where $\boldsymbol{w}$ is a diagonal local weight matrix. Application of Eq. (8) leads to Algorithm 2. Note that $\boldsymbol{w} = \boldsymbol{I}$ recovers the classical additive Schwarz method.

Table 1: Weight functions (WF) and related shape functions

WF	shape function	method
w_1	$\hat{\phi}_1 = x$	linear
w_3	$\hat{\phi}_3 = (3x - x^3)/2$	cubic
w_5	$\hat{\phi}_5 = (15x - 10x^3 + 3x^5)/8$	quintic
w_7	$\hat{\phi}_7 = (35x - 35x^3 + 21x^5 - 5x^7)/16$	7^{th} order
w_t	$\hat{\phi}_t = \mathrm{sgn}(x)$	top hat
w_a	$\hat{\phi}_a = 0$	arithmetic mean

However, Lottes and Fischer [19] found that weighting is essential to effective Schwarz smoothing in multigrid methods and proposed to use the arithmetic mean of local corrections. In our notation, this corresponds to choosing $\boldsymbol{w} = \boldsymbol{R}_s \boldsymbol{C}^+ \boldsymbol{R}_s^t$, where $\boldsymbol{C}^+$ is the pseudoinverse of the counting matrix $\boldsymbol{C} = \sum_s \boldsymbol{R}_s^t \boldsymbol{R}_s$. As an alternative we consider a more flexible approach which elevates the weights gradually from zero at the border to one in that part of Ω_s which is not covered by other subdomains. Due to the regular shape of Ω_s the weights can be cast in the tensor product form $\boldsymbol{w} = \tilde{\boldsymbol{w}} \otimes \tilde{\boldsymbol{w}}$. The one-dimensional weight distribution $\tilde{\boldsymbol{w}}$ is generated from the continuous weighting function

$$w_i(\xi) = \frac{1}{2}\left[1 + \phi_i\left(\frac{\xi+1}{\delta_\mathrm{o}}\right) - \phi_i\left(\frac{\xi-1}{\delta_\mathrm{o}}\right)\right], \tag{9}$$

where ξ is the 1D standard coordinate extended beyond $\hat{\Omega}$ and ϕ_i is a weakly monotonic increasing shape function. In particular we consider the shape functions ϕ_i with $i \in \{1, 3, 5, \dots\}$ defined as

$$\phi_i(x) = \begin{cases} \hat{\phi}_i(x) & x \in \hat{\Omega} \\ \mathrm{sgn}(x) & \text{else} \end{cases} \tag{10}$$

where $\hat{\phi}_i$ is a polynomial of degree i satisfying the conditions

$$\hat{\phi}_i(\pm 1) = \pm 1 \tag{11a}$$

$$\frac{\mathrm{d}^k \hat{\phi}_i}{\mathrm{d}\, x^k}(\pm 1) = 0, \quad 0 < k \le (i-1)/2\,. \tag{11b}$$

The $\hat{\phi}_i$ are strictly monotonic in $(-1, 1)$ and generate a smooth transition of the weight function in the overlap zone, as exemplified in Fig. 2 for the quintic case. By increasing the polynomial degree the shape function converges toward the sign function, which translates into a top hat weighting function. Finally, we remark that omitting the shape function in Eq. (9) yields the arithmetic mean. For reference, Table 1 summarizes all weight functions used in the numerical experiments.

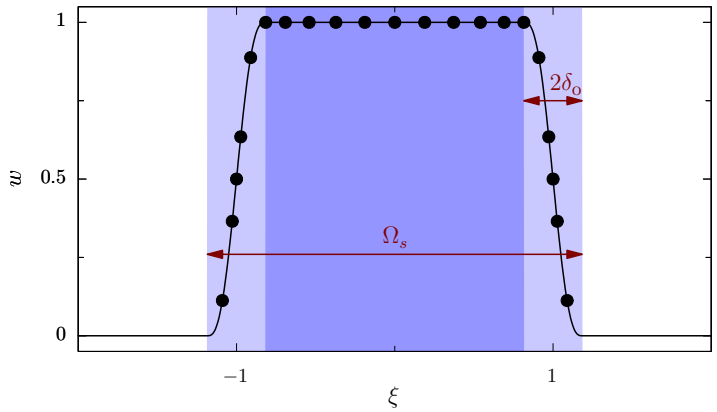

Fig. 2: One-dimensional weight distribution for elements of order $p = 16$ with an overlap of $n_o = 2$ points using a quintic shape function. The core region and the overlap zone of the subdomain are shaded in dark and light blue, respectively. Filled circles indicate the node positions.

3.2 Multigrid

For MG we define a series of polynomial levels $\{p_l\}$ with $p_l = 2^l$ increasing from 1 at $l = 0$ to p at top level L. Correspondingly, $\boldsymbol{u}_l$ denotes the global coefficients and $\boldsymbol{A}_l$ the system matrix on level l. On the top level we have $\boldsymbol{u}_L = \boldsymbol{u}$ and $\boldsymbol{A}_L = \boldsymbol{A}$, whereas on lower levels $\boldsymbol{u}_l$ is the defect correction and $\boldsymbol{A}_l$ the counterpart of $\boldsymbol{A}$ obtained with elements of order p_l. For transferring the correction from level $l-1$ to level l we use the embedded interpolation operator $\boldsymbol{\mathcal{I}}_l$, and for restricting the residual its transpose. These ingredients allow to build a multigrid V-cycle, which is summarized in Algorithm 3. Both, the multiplicative and the weighted additive Schwarz method stated in Algorithm 1 and 2, respectively, can serve as the SMOOTHER. The number of pre- and post-smoothing steps, $n_{s1,l}$ and $n_{s2,l}$, can differ from level to level to allow variable V cycles [2]. Line 8 of Algorithm 3 defines the coarse grid solution by means of the pseudoinverse $\boldsymbol{A}_0^+$. In our implementation the coarse problem is solved using the conjugate gradient method. To achieve convergence in spite of singularity, the right side is projected to the null space of $\boldsymbol{A}_0$, as proposed by Kaasschieter [14].

3.3 Preconditioned conjugate gradients

For enhancing robustness and efficiency multigrid methods can be accelerated by Krylov subspace methods [27]. In the present case, with symmetric system matrices on all grid levels, one would favor preconditioned conjugate gradients. Unfortunately, weighted additive Schwarz and multiplicative Schwarz with un-

Algorithm 3 Multigrid V-cycle.

```
1: function MultigridCycle(u, f, n_s)
2:     u_L ← u
3:     f_L ← f
4:     for l = L, 1 step −1 do
5:         u_l ← Smoother(u_l, f_l, n_{s1,l})                ▷ Pre-smoothing
6:         f_{l−1} ← I_l^t (f_l − A_l u_l)                   ▷ Residual restriction
7:     end for
8:     u_0 ← A_0^+ f_0                                       ▷ Coarse grid solution
9:     for l = 1, L do
10:        u_l ← u_l + I_l u_{l−1}                           ▷ Correction prolongation
11:        u_l ← Smoother(u_l, f_l, n_{s2,l})                ▷ Post-smoothing
12:    end for
13:    return u ← u_L
14: end function
```

Algorithm 4 Inexact multigrid preconditioned conjugate gradients.

```
1: function MGCG(u, f, n_s, i_max, r_max)
2:     r_old ← 0
3:     r ← f − Au
4:     p ← MultigridCycle(0, r, n_s)
5:     δ ← p^t r
6:     for i = 1, i_max do
7:         q ← Ap
8:         α ← δ/(p^t q)
9:         u ← u + αp
10:        r ← r − αq
11:        if ‖r‖ ≤ r_max exit
12:        z ← MultigridCycle(0, r, n_s)
13:        β ← q^t (r − r_old)/δ
14:        p ← z + βp
15:        δ ← z^t r
16:        r_old ← r
17:    end for
18:    return u
19: end function
```

even iteration count are both non symmetric and hence affect the symmetry of MG as well. As a remedy, Lottes and Fischer [19] symmetrized their weighting method or, alternatively, used GMRES. According to Loisel et al. [18], however, symmetrization can worsen the efficiency of the method and is therefore not recommended. Moreover, several generalizations of standard CG have been developed that seem to be well suited for the present problem [8, 23]. Here we adopt the inexact preconditioned conjugate gradients of Golub and Ye [8]. Algorithm 4 summarizes the resulting MGCG method. Note that, as before with the coarse problem, the right side $\boldsymbol{f}$ must be in the null space of $\boldsymbol{A}$ if the system is singular.

4 Results

Numerical tests were performed to assess the influence of weighting, overlap and cycling strategy on the computational efficiency and robustness of MG and MGCG. The methods were implemented in Fortran and compiled using the GNU compiler collection 6.0 with -O3.

All results reported here are based on the test case of Lottes and Fischer [6, 19] with $f = 2\pi^2 \sin(\pi x)\sin(\pi y)$ and starting from a random initial guess confined to the interval [0,1]. The convergence speed is evaluated using the number of cycles n_{10} needed to reduce the norm of the residual by a factor of 10^{10} and the average logarithmic convergence rate according to Varga [28]

$$\bar{r} = \frac{1}{n} \log_{10} \frac{\|\boldsymbol{r}^{(0)}\|}{\|\boldsymbol{r}^{(n)}\|} , \tag{12}$$

where $\boldsymbol{r}^{(n)}$ is the Euclidean norm of the residual vector after the nth cycle. Note that n_{10} is nearest integer greater than or equal to $10/\bar{r}$.

As an efficiency measure we define the approximate number of operator applications required for reducing the residual by a factor of 10^k,

$$\bar{\omega}_k = \frac{k}{\bar{r}} \frac{W_{\text{cyc}}}{W_{\text{op}}} , \tag{13}$$

where W_{cyc} is the cost for one cycle and W_{op} the cost for one application of the system matrix $\boldsymbol{A}$. Exploiting sum factorization [5, 16], W_{op} can be estimated as $2n_{\text{P}}^3 n_{\text{E}}$, where $n_{\text{P}} = p + 1$. According to Sec. 3.1, the cost of one Schwarz iteration is approximately $2(n_{\text{P}} + 2n_{\text{O}})^3 n_{\text{E}}$. Assuming a maximum relative overlap of $n_{\text{O}}/n_{\text{P}}$ this yields the estimate

$$W_{\text{cyc}} = \left[4\Big(1 + 2\frac{n_{\text{O}}}{n_{\text{P}}}\Big)^3 c_{\text{S}} n_{\text{S}} + 2c_{\text{S}} + c_{\text{CG}} \right] n_{\text{P}}^3 n_{\text{E}} , \tag{14}$$

where n_{S} is the number of pre- and post-smoothing steps on the finest level, $c_{\text{S}} = 4/3$ for the classical V-cycle and $c_{\text{S}} = 2$ for a variable V-cycle doubling the number of smoothing steps when changing to the next lower level, and $c_{\text{CG}} = 2$ is the extra cost for conjugate gradients with MGCG.

4.1 Weighting and overlap

For this study we adopted the test case of Lottes and Fischer [19] featuring a square domain with length $l_x = l_y = 2$, which is uniformly subdivided in 8×8 square elements with order p ranging from 4 to 32. In the first test series we set the overlap to $n_{\text{O}} = 1$ on all levels $l > 0$. Table 2 shows the measured convergence rates for MG with one pre-smoothing. Column "w_a" corresponds to the weighted additive Schwarz method of Lottes and Fischer [19] using the arithmetic mean in overlap areas. Compared to [19] our results agree well for $p = 4$, but show a faster convergence for higher orders. This could be attributed to

Table 2: Convergence rates $\bar{r}$ for MG with additive Schwarz smoothing using one pre- and no post-smoothing steps, $n_\mathrm{E} = 8 \times 8$ elements and a fixed overlap of $n_\mathrm{O} = 1$. The weighting methods are referred to as defined in Tab. 1. Results for the multiplicative smoother (mult) are included for comparison.

p	w_a	w_1	w_3	w_5	w_7	w_t	mult
4	0.66	0.86	1.01	1.17	1.25	0.72	1.01
8	0.40	0.83	1.17	1.29	1.23	0.52	1.29
16	0.34	0.80	0.84	0.84	0.84	0.42	1.26
32	0.32	0.43	0.43	0.43	0.43	0.38	0.76

Table 3: Convergence rates for a level dependent overlap of $n_{\mathrm{O},l} = \lfloor p_l/8 \rfloor$. For caption see Tab 2.

p	w_a	w_1	w_3	w_5	w_7	w_t	mult
4	0.63	0.91	0.98	0.96	0.79	0.31	1.03
8	0.40	0.75	1.06	1.28	1.28	0.64	1.30
16	0.51	1.07	1.36	1.28	1.12	0.53	1.40
32	0.71	1.39	1.48	1.50	1.51	0.19	1.56

using periodic instead of Dirichlet boundary conditions. We remark, however, that the iteration counts obtained with the two-grid method (not shown here) agree excellently with those reported in [19].

The remaining columns in Tab. 2 display the convergence rates for additive Schwarz smoothing with the gradual weighting introduced in Sec. 3.1 and multiplicative Schwarz. In comparison to the arithmetic mean, weighting based on a smooth – cubic, quintic or 7^{th} order – shape function roughly doubled the convergence rate for orders 4, 8 and 16, while linear and top hat weighting yielded a lower, but still remarkable improvement. As it was to be expected, the multiplicative Schwarz smoother attained the fastest convergence. At $p = 32$ all methods suffer a serious performance degradation, except for arithmetically weighted Schwarz, which nonetheless remains the slowest.

Inspired by these observations, several tests were run with overlaps depending on the polynomial gegree on each mesh level. Table 3 shows the convergence rates for the case $n_{\mathrm{O},l} = \lfloor p_l/8 \rfloor$. Note that this choice implies $n_\mathrm{O} = 0$ for degrees less than 8, while reaching $n_\mathrm{O} = 4$ with $p = 32$. As a consequence, the convergence rates for $p = 4$ are slightly lower than with $n_\mathrm{O} = 1$, except for multiplicative Schwarz. For $p \geq 16$ the increased overlap yields a considerable speedup. This improvement is most pronounced for cubic and quintic weighting, which come remarkably close to multiplicative Schwarz.

As a résumé of the first study we conclude that 1) gradual weighting with a smooth shape function yields a decisive improvement over arithmetic weighting, and 2) increasing the overlap with growing p is crucial for robustness.

4.2 Robustness and efficiency

Next we investigate robustness with respect to the mesh size and aspect ratio. First, MG with one pre-smoothing is applied on uniform meshes consisting of 4^2 to 1024^2 elements with p ranging from 4 to 32 and up to four million unknowns. Table 4 compiles the results for quintic weighted and multiplicative Schwarz smoothers with overlap $n_{\mathrm{o},l} = \max(1, \lfloor p_l/8 \rfloor)$. Except of coarse quadrangulations, where periodicity can induce interference effects, the convergence characteristics applications are virtually independent of the number of elements n_{E}. The convergence rate $\bar{r}$ shows a moderate growth for increasing order p and is similar for both smoothers, with a slight advantage for the weighted additive Schwarz method. As a consequence, the equivalent number of operator applications required for reducing the residual by an order of magnitude drops to almost one third when increasing p from 4 to 32 and, thus, mitigates the higher operator cost per DOF.

In the second test we fixed the mesh to 16×16 elements of order $p = 16$, but increased the aspect ratio $AR = \Delta x/\Delta y$ by enlarging the domain into the x direction. Table 5 reports the results for MG and MGCG using additive weighted Schwarz with w_5, $n_{\mathrm{o},l} = \max(1, \lfloor p_l/8 \rfloor)$ and one pre-smoothing step. As expected, the stand-alone MG performs well for small aspect ratios, but degrades for $AR > 2$. MGCG is slightly less efficient than MG for $AR \leq 2$, but proves more robust at higher aspect ratios. At $AR = 8$ it converges approximately twice as fast as MG.

While these observations hold almost uniformly for all orders p considered, it remains to investigate the impact of solver parameters such as smoothing steps and overlap. Figure 3 presents selected results of the corresponding study for $p = 16$ and aspect ratios $AR = 1$ to 16. In particular we considered several variants of MGCG(1,1), each applying one pre- and one post-smoothing. In one case, indicated by "var", we employed a variable V-cycle in which the number of smoothing steps doubles with each coarser level, i.e. $n_{\mathrm{s}1,l} = n_{\mathrm{s}2,l} = 2^{L-l}$. The study included quintic weighted additive ("add, w_5") as well as multiplicative ("mult") Schwarz smoothers with degree dependent overlap, $n_{\mathrm{o},l} = \max(1, \lfloor p_l/8 \rfloor)$. Additionally we tested multiplicative Schwarz with $n_{\mathrm{o}} = 0$ and $n_{\mathrm{s}1} = n_{\mathrm{s}2} = 2$, which corresponds to the method of Janssen and Kanschat [13], and the arithmetically averaged additive Schwarz smoother of Lottes and Fischer [19] using a constant overlap of $n_{\mathrm{o}} = 1$. Figure 3a depicts the achieved convergence rates. Compared to the case of only one smoothing, the additional post-smoothing raises $\bar{r}$ by a factor between 1.5 and 2, which is well in the expected range. Switching to multiplicative Schwarz yields an even higher gain for higher aspect ratios. A similar effect is achieved using additive Schwarz with the variable V cycle. MGCG(2,2) with zero overlap attains a convergence rate similar to MGCG(1,0) with degree-dependent overlap. The arithmetically averaged Schwarz method with two smoothing steps falls about two thirds behind the quintic weighted method with only one smoothing for $AR = 1$, but gains a slight advantage over the latter for higher aspect ratios.

Table 4: Robustness with respect to problem size: MG using additive and multiplicative Schwarz smoothers with overlap $n_{\mathrm{O},l} = \max(1, \lfloor p_l/8 \rfloor)$.

		MG(1,0), add, w_5			MG(1,0), mult		
p	$\sqrt{n_\mathrm{E}}$	$\bar{r}$	n_{10}	$\bar{\omega}_1$	$\bar{r}$	n_{10}	$\bar{\omega}_1$
4	32	1.17	9	9.2	0.87	12	12.4
	64	1.17	9	9.3	0.86	12	12.6
	128	1.17	9	9.3	0.85	12	12.7
	256	1.17	9	9.3	0.85	12	12.7
8	16	1.30	8	5.4	1.28	8	5.5
	32	1.29	8	5.4	1.26	8	5.5
	64	1.29	8	5.4	1.26	8	5.5
	128	1.28	8	5.4	1.26	8	5.5
16	8	1.33	8	5.1	1.44	7	4.7
	16	1.37	8	4.9	1.42	8	4.8
	32	1.36	8	5.0	1.46	7	4.6
	64	1.36	8	5.0	1.46	7	4.6
32	4	1.90	6	3.5	1.65	7	4.0
	8	1.58	7	4.2	1.59	7	4.2
	16	1.87	6	3.6	1.63	7	4.1
	32	1.93	6	3.4	1.64	7	4.0
	64	1.93	6	3.4	1.65	7	4.0

Table 5: Robustness with respect to aspect ratio: MG versus MGCG using additive Schwarz with w_5 and overlap $n_{\mathrm{O},l} = \max(1, \lfloor p_l/8 \rfloor)$.

		MG(1,0), add, w_5			MGCG(1,0), add, w_5		
p	AR	$\bar{r}$	n_{10}	$\bar{\omega}_1$	$\bar{r}$	n_{10}	$\bar{\omega}_1$
4	1	1.17	9	9.2	1.30	8	9.3
	2	0.99	11	11.0	1.10	10	11.0
	4	0.39	26	27.8	0.59	17	20.5
	8	0.12	85	91.2	0.28	36	42.5
8	1	1.30	8	5.4	1.33	8	6.1
	2	0.86	12	8.2	1.03	8	7.9
	4	0.43	24	16.3	0.65	16	12.5
	8	0.16	63	43.9	0.34	30	23.8
16	1	1.37	8	4.9	1.55	7	5.1
	2	0.95	11	7.1	1.14	9	6.9
	4	0.50	20	13.5	0.72	14	10.8
	8	0.17	59	39.9	0.39	26	20.1
32	1	1.87	6	3.6	2.01	5	3.8
	2	1.23	9	5.4	1.42	8	5.4
	4	0.65	16	10.3	0.83	12	9.2
	8	0.22	46	30.4	0.44	23	17.5

As it does not account for the cost, the convergence rate is of limited value when comparing methods of different computational complexity. A better measure is the equivalent number of operator applications required for reducing the residual by one order of magnitude, $\bar{\omega}_1$, which is shown in Fig. 3b. In this metric, the multiplicative MGCG(1,1) with degree dependent overlap performs best, especially for higher aspect ratios. It is followed by its quintic weighted counterpart, which is at level for $AR \leq 2$, but needs ca 34 instead of 26 operator applications for $AR = 16$. The comparison also reveals that the benefit of the variable V-cycle is lost due to the higher computational complexity. Generally, the influence of smoothing and overlap parameters lessens with increasing aspect ratio (exempting the case of $n_\mathrm{o} = 0$), which indicates that the role of the conjugate gradient method gets more important.

Finally, Figure 3c depicts the runtimes measured on a 3.1 GHz Intel Core i7-5557U CPU. Notably, the quintic weighted MGCG(1,1) attained the best performance despite the higher operation count in comparison to MGCG(1,1) with multiplicative Schwarz. This is because the additive Schwarz method evaluates the residual for all elements at once, yielding a single, highly efficient BLAS 3 operation. In contrast, multiplicative Schwarz requires a series of local residual updates, which is harder to optimize. Consistently, the multiplicative MGCG(2,2) with $n_\mathrm{o} = 0$ remains the least efficient method for all aspect ratios. Compared to arithmetically weighted MGCG(1,1) with $n_\mathrm{o} = 1$, the quintic weighted method with $n_{\mathrm{o},l} = \max(1, \lfloor p_l/8 \rfloor)$ succeeds twice as fast for $AR = 1$ and still gains 23% at $AR = 16$. Though other choices may exist that yield even better performance, the study documents that the method is not too sensitive to parameter variations such that only minor improvements can be expected.

5 Conclusions

We have developed a non-uniformly weighted additive Schwarz method acting as the smoother in multigrid solvers for the spectral element discretization of the Poisson equation in $\mathbb{R}^2$. The method can be seen as a generalization of the arithmetically weighted Schwarz method of Lottes and Fischer [19] and was inspired from weighting techniques devised in [10]. In each step, it determines a solution for a subdomain corresponding to an extended element region. These local solutions are blended according to a polynomial shape function which features a smooth transition from zero at the border toward one in the core of the subdomain. As an alternative we considered a multiplicative Schwarz method with no weighting required. Both Schwarz methods were integrated in a polynomial multigrid method which, in turn, was embedded in the inexact preconditioned CG method of Golub and Ye [8].

The performance of these methods was assessed in a series of numerical experiments with ansatz orders p ranging from 4 to 32 and up to $n_\mathrm{E} = 256^2$ elements of aspect ratios AR from 1 to 16. For unit-aspect ratio elements the proposed weighting improved the logarithmic MG convergence rate by a factor of 1.5 to 3 in comparison to the original method [19]. The study

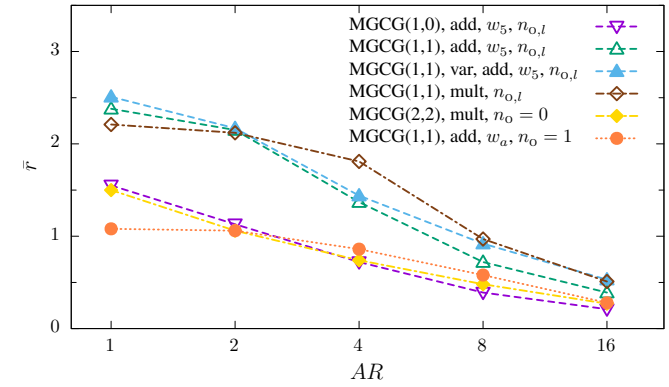

(a) Average logarithmic convergence rate

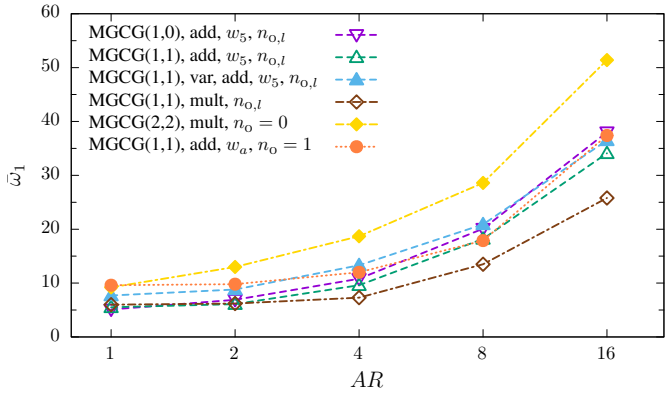

(b) Operator applications required for 10^1 residual reduction

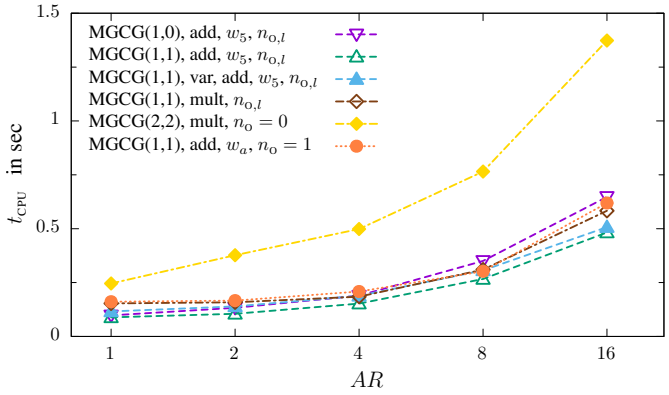

(c) Solver runtime for 10^{10} residual reduction

Fig. 3: Performance of MGCG for different aspect ratios. All cases use 16×16 elements of order $p = 16$ and $n_{o,l} = \max(1, \lfloor p_l/8 \rfloor)$, if not specified otherwise.

indicates that for robustness the subdomain overlap has to be bounded, i.e., the number n_{O} of node layers adopted from neighbor elements must grow with increasing order. Thus, with MG, the number of layers varies from level to level. In our tests $n_{\mathrm{O},l} = \max(1, \lfloor p_l/8 \rfloor)$ proved to be a reasonable choice. The resulting multigrid method is robust with respect to the mesh size (i.e. p and n_{E}), but degrades with increasing aspect ratio. This behavior can be mitigated by Krylov subspace acceleration: Using MG as a preconditioner for the inexact conjugate gradient method [8] improves the convergence rate for higher aspect ratios considerably. For example, with $AR = 8$, the accelerated method consistently needed less than half the iteration count of MG for all ansatz orders.

For moderate aspect ratios the improved additive Schwarz method the competes easily with multiplicative Schwarz. Using one pre- and post-smoothing, the preconditioned MG solvers both require between 5 and 6 operator applications for reducing the residual by one order of magnitude. For $AR > 4$ the multiplicative method converges somewhat faster. However, this advantage is more than compensated by the higher computational efficiency of the additive method.

We remark that the methods presented here are less efficient in 2D than the condensed multigrid method presented in [10]. Nonetheless they may provide a better starting point in the three-dimensional case, as the fast diagonalization techniques crucial for optimal complexity do not easily extend to the condensed operators.

References

1. P. Bastian, M. Blatt, and R. Scheichl. Algebraic multigrid for discontinuous Galerkin discretizations of heterogeneous elliptic problems. *Numer. Linear Algebra Appl.*, 19(2):367–388, 2012. ISSN 1099-1506.
2. J. Bramble. *Multigrid methods.* Pitman Res. Notes Math. Ser. 294. Longman Scientific & Technical, Harlow, UK, 1995.
3. S. C. Brenner and J. Zhao. Convergence of multigrid algorithms for interior penalty methods. *Appl. Num. Anal. Comp. Math.*, 2(1):3–18, 2005.
4. W. Couzy and M. O. Deville. A fast Schur complement method for the spectral element discretization of the incompressible Navier-Stokes equations. *J. Comput. Phys.*, 116:135–142, January 1995. ISSN 0021-9991.
5. M. O. Deville, P. F. Fischer, and E. H. Mund. *High-Order Methods for Incompressible Fluid Flow*, volume 1. Cambridge University Press, 2002.
6. P. F. Fischer. Personal communication, September 2015.
7. P. F. Fischer and J. W. Lottes. Hybrid schwarz-multigrid methods for the spectral element method: Extensions to navier-stokes. In *Domain Decomposition Methods in Science and Engineering Series*, pages 35–49. Springer, 2004.

8. G. H. Golub and Q. Ye. Inexact preconditioned conjugate gradient method with inner-outer iteration. *SIAM Journal on Scientific Computing*, 21(4): 1305–1320, Dec. 1999. ISSN 10648275.
9. J. L. Guermond, P. Minev, and J. Shen. An overview of projection methods for incompressible flows. *Comput. Methods Appl. Mech. Eng.*, 195:6011–6045, 2006.
10. L. Haupt, J. Stiller, and W. Nagel. A fast spectral element solver combining static condensation and multigrid techniques. *J. Comput. Phys.*, 255: 384–395, 2013.
11. W. Heinrichs. Line relaxation for spectral multigrid methods. *J. Comput. Phys.*, 77:166–182, 1988.
12. I. Huismann, L. Haupt, J. Stiller, and J. Fröhlich. Sum factorization of the static condensed helmholtz equation in a three-dimensional spectral element discretization. *PAMM Proc. Appl. Math.*, 14(1):969–970, 2014.
13. B. Janssen and G. Kanschat. Adaptive multilevel methods with local smoothing for H^1- and H^{curl}-conforming high order finite element methods. *SIAM J. Sci. Comput.*, 33(4):2095–2114, 2011.
14. E. F. Kaasschieter. Preconditioned conjugate gradients for solving singular systems. *J. Comput. Appl. Math.*, 24:265–275, 1988.
15. G. Kanschat. Multilevel methods for discontinuous Galerkin FEM on locally refined meshes. *Computers and Structures*, 82:2437–2445, 2004.
16. G. E. Karniadakis and S. J. Sherwin. *Spectral/hp Element Methods for Computational Fluid Dynamics*. Oxford University Press, 2nd edition, 2005.
17. J. K. Kraus and S. K. Tomar. A multilevel method for discontinuous galerkin approximation of three-dimensional anisotropic elliptic problems. *Numer. Linear Algebra Appl.*, 15(5):417–438, 2008.
18. S. Loisel, R. Nabben, and D. Szyld. On hybrid multigrid-schwarz algorithms. *J. Sci. Comput.*, 36(2):165–175, 2008. ISSN 0885-7474.
19. J. W. Lottes and P. F. Fischer. Hybrid multigrid/Schwarz algorithms for the spectral element method. *J. Sci. Comput.*, 24:45–78, 2005. ISSN 0885-7474.
20. R. E. Lynch, J. R. Rice, and D. H. Thomas. Direct solution of partial difference equations by tensor product methods. *Numer. Math.*, 6:185–199, 1964. ISSN 0029-599X.
21. Y. Maday and R. Munoz. Spectral element multigrid. II. Theoretical justification. *J. Sci. Comput.*, 3:323–353, 1988. ISSN 0885-7474.
22. W. F. Mitchell. The hp-multigrid method applied to hp-adaptive refinement of triangular grids. *Numer. Linear Algebr.*, 17:211–228, 2010.
23. Y. Notay. Flexible conjugate gradients. *SIAM Journal on Scientific Computing*, 22(4):1444–1460, 2000.
24. L. Olson. Algebraic multigrid preconditioning of high-order spectral elements for elliptic problems on a simplicial mesh. *SIAM J. Sci. Comput.*, 29(5):2189–2209, 2007. ISSN 1064-8275.
25. R. Pasquetti and F. Rapetti. p-multigrid method for Fekete-Gauss spectral element approximations of elliptic problems. *Commun. Comput. Phys.*, 5

(5):667–682, February 2009.
26. E. Rønquist and A. Patera. Spectral element multigrid. I. Formulation and numerical results. *J. Sci. Comput.*, 2:389–406, 1987. ISSN 0885-7474.
27. U. Trottenberg, C. W. Oosterlee, and A. Schüller. *Multigrid*. Academic Press, 2000.
28. R. S. Varga. *Matrix Iterative Analysis*. Springer Ser. Comput. Math. 27. Springer-Verlag, Berlin, 2nd edition, 2000.
29. Z. J. Wang, K. Fidkowski, R. Abgrall, F. Bassi, D. Caraeni, A. Cary, H. Deconinck, R. Hartmann, K. Hillewaert, H. Huynh, N. Kroll, G. May, P.-O. Persson, B. van Leer, and M. Visbal. High-order cfd methods: current status and perspective. *International Journal for Numerical Methods in Fluids*, 72(8):811–845, 2013. ISSN 1097-0363.